

RESOURCE ADMISSION CONTROL IN AN ACCESS NETWORK

The present invention relates to a method to guarantee for a service a bandwidth across an access network with a quality of service, as described in the preamble of claim 1, and to an access network wherein the method is applied, as described in the preamble of claim 7.

Such a method is already known in the art, e.g. from the document entitled '*Traffic Management Specification*', published by the ATM forum technical committee on March 1999, version 4.1, reference AF-TM-0121.000. Therein, a method to achieve performance objectives in an Asynchronous Transfer Mode (ATM) network and to protect the network from traffic congestion is described.

An access network is that part of a data communications network that connects the subscribers to the core network. The access network may overlay upon an established network such as the Public Switched Telephone Network (PSTN), typically upon the last few miles of the PSTN to serve thinly scattered residential subscribers. The trade-off is to get the highest traffic concentration ratios across the access network, i.e. less equipment and less leased lines, while guaranteeing the best network performances to the subscribers.

ATM is a candidate technology for the transport layer between the subscriber and the core network. ATM is a connection-oriented cell-based transport technique and is intended to support a wide variety of services and applications. At the subscriber premises, a Data Communication Equipment (DCE) chops the upstream traffic from a Data Terminal Equipment (DTE) into data blocks, encapsulates the data blocks into ATM cells and transmits the ATM cells over the User Network Interface (UNI). The ATM cells are transported over a virtual connection, which is set up across the access network between the DCE and an edge node of the access network coupled to the core network. The edge node extracts the upstream traffic from the ATM cells and transmits it over the core network. In the reverse direction, the edge node encapsulates the

downstream traffic from the core network into ATM cells and transmits them over the appropriate virtual connection. The edge node could also stand for:

- a network element of the access network that is coupled to the PSTN,
5 such as a voice gateway,
- a network element of the access network that houses a service provider or that is coupled via a Local Area Network (LAN) to a service provider, or any alternative as known to the skilled person.

ATM defines a two-level hierarchical connection scheme:

- 10 - the Virtual Channel (VC) is a connection between two end systems for a specific service or application,
- the Virtual Path (VP) is a connection between two network nodes that transports a large number of simultaneous VCs.

A connection can be initiated by any of the two end systems
15 involved in the connection or can be provisioned by a third party, such as a network manager. In the latter case, the connection is referred to as Permanent Virtual Channel (PVC) or Permanent Virtual Path (PVP).

The connection-oriented mode of operation allows the network to check whether enough resources are available for a connection and then to
20 commit some Quality of Service (QoS) for that connection. The bandwidth and the QoS required for a connection are negotiated at connection establishment time. The network can reserve therefrom the appropriate resources for this connection. The network guarantees that the congestion of the network resources will only occur with a very low pre-determined
25 probability.

An object of the known method is to meet bandwidth and QoS requirements across an ATM network. Another object is to promote the efficient use of the network resources.

In a preliminary step of the known method, a set of parameters is
30 defined for describing the bandwidth and the QoS requirements.

The bandwidth requirements are expressed by means of traffic descriptors featuring the way a source may transmit cells over the network. The Peak Cell Rate (PCR) defines the maximum cell rate a source may reach. In case of fluctuating cell rate (e.g., because of a video compression
5 algorithm), additional descriptors might be specified to allow the network to allocate resources more efficiently.

The quality of service of a connection relates to the cell loss, the delay and the delay variation (also called jitter) incurred by the cells transferred over the connection.

10 Thereupon, appropriately differentiated service categories are defined. Whenever a connection is set up, the service category to which that connection relates and the traffic and QoS parameters relevant to that service category are specified.

The ATM service categories are:

- 15 - Unspecified Bit Rate (UBR) service category,
- Constant Bit Rate (CBR) service category,
- real time – Variable Bit Rate (rt-VBR) service category,
- non-real time – Variable Bit Rate (nrt-VBR) service category,
- Available Bit Rate (ABR) service category,
20 - Guaranteed Frame Rate (GFR) service category.

The UBR service category is intended for non-real time best effort applications, such as internet browsing, file transfer or e-mail.

The CBR service category is used by connections that request a static amount of bandwidth that is continuously available during the
25 connection lifetime. This amount of bandwidth is characterized by the PCR value. The source may emit cells at or below the negotiated PCR (and may also even be silent) at any time and for any duration and the QoS commitments still pertain. CBR service is intended for real time applications requiring tightly constrained delay, delay variation and cell
30 loss ratio, such as video streaming or Voice over IP (VoIP).

Refer to the above mentioned specification for further details about the remaining service categories and the set of parameters relevant to each category.

5 In a further step of the known method, mechanisms are specified for managing and controlling traffic and congestion in ATM networks. These mechanisms are:

- Connection Admission Control (CAC), which is defined as the set of actions taken by the network during the connection set up phase to determine whether a connection request can be progressed or should be
10 rejected according to the network resources that remain available for that connection,
- Usage Parameter Control (UPC), which is defined as the set of actions taken by the network to monitor and, possibly, to enforce the traffic contract at the UNI,
- 15 - Network Parameter Control (NPC), which is defined as the set of actions taken by the network to monitor and, possibly, to enforce the traffic contract at the Network Node Interface (NNI),
- Feedback control, which is defined as the set of actions taken by the network and by the end-systems to adjust the traffic submitted on a ATM
20 connection according to the state of the network elements along that connection,
- Cell Loss Priority (CLP) control, which allows the network to discard some lower priority cells to protect, as far as possible, the QoS objectives of higher priority cells,
- 25 - Traffic Shaping, used to achieve a desired modification to the traffic characteristics of a connection.

The main objective of the UPC and NPC policers is to protect the network resources from malicious, as well as unintentional misbehaviour, which can affect the QoS of other established connections. This protection is
30 achieved by detecting violations of the negotiated traffic parameters and by taking appropriate actions, such as discarding cells.

Most of the ATM based access networks make use of PVCs to connect a subscriber's DCE to an edge node of the access network. As a matter of fact, the installed base of subscriber's DCEs are for the most part unable to set up an ATM connection on their own. They miss important
5 ATM features (no signalling channel, etc) as well as the necessary pieces of information (no ATM addressing data, etc).

A subscriber is provisioned with as many PVCs as there are classes of service to support for that subscriber. Each PVC is assigned a suitable service category in accordance with the known method.

10 A shortcoming of the know method is the very loose integration between those PVCs overlaying upon the access network and adapted to transport an amount of traffic across the access network with a QoS, and the higher communication layers submitting traffic thereto. The traffic submitted over a virtual connection does not necessarily need to conform to
15 the ATM traffic contract of that virtual connection. If the real traffic load goes beyond the provisioned bandwidth, a UPC or NPC policer will start discarding ATM cells irrespective of the service to which the ATM cells relate. Consequently, delivering a new service over a connection without any further check may alter the integrity of existing services delivered over that
20 connection.

An object of the present invention is to improve service integrity by controlling the right for a service instance to use the network resources that have been provisioned for its class of service.

According to the invention, this object is achieved by the method
25 defined in claim 1 and by the access network defined in claim 7.

Whenever a subscriber requests a service that requires a pre-determined bandwidth across an access network with a pre-determined QoS, the service is granted or denied according to the resources of the access network that remain available to deliver said service to said subscriber with
30 said quality of service.

First of all, the virtual connection(s) that has been provisioned across the access network for that class of service and that subscriber is identified. The identification step comprises the steps of:

- 5 - identifying the edge node of the access network through which said service is delivered to said subscriber,
- translating the data transfer capabilities, as required by said service with respect to the data communication network, into data transfer capabilities, as required by said service with respect to the access network,
- 10 - retrieving the configuration of the access network,
- identifying which one(s) of the virtual connections that have been provisioned between said subscriber and the so-identified edge node supports the so-translated data transfer capabilities,

Next, one checks whether the so-identified virtual connection(s)
15 can convey the required bandwidth with the required QoS while maintaining the QoS agreed with the services already sharing that virtual connection,

According to the check outcome, said subscriber is granted or denied said service.

20 By checking upon service request whether an access network can accommodate an additional amount of traffic, one can further overcome the over provisioning of the network resources. In the known method, the bandwidth and the service category of a PVC is determined from a service level agreement between the subscriber and the operator. A PVC is
25 allocated the necessary network resources to convey at any time and for any duration the traffic volume agreed with the subscriber and committed upon by the operator, even though that PVC might not be active, i.e. might not convey any traffic. An alternative would be to define only UBR PVCs, as prevailing in most of the access networks. However, the UBR PVCs are
30 intended for best effort traffic and thus are helpless for network applications

that requires QoS commitments for the data transferred across the access network.

Several embodiments can be thought off that realizes the optimization of the network resources in an access network across which
5 virtual connections have been provisioned to meet traffic and QoS requirements, while still improving the integrity of the services delivered over those virtual connections.

Three of them are defined in the claims 2 and 8, 3 and 9, 4 and 10 respectively.

10 In the first embodiment (claims 2 and 8), one adapts the capacity of the so-identified virtual connection for it to convey the required bandwidth with the required QoS, provided there are enough resources available in the access network to do so, or alternatively provided the CAC of the network elements across which the virtual connection is established
15 grants the permission to do so. One can either provision a suitable bandwidth for that virtual connection or set the policer parameters to allow more cells to cross the UNI and the NNI. By fitting closer to the real traffic demand, one can transport more traffic across the access network than one could if the virtual connections were left as provisioned initially.

20 In the second and third embodiments (claims 3 and 9, 4 and 10 respectively), a path is provisioned across the access network, over which a number of virtual connections are multiplexed. The path is provisioned between the node whereto the subscribers are connected and the edge node, or between any intermediate nodes.

25 A path stands for any amount of bandwidth that connects a network element of the access network to another. It could simply be a physical connection link or a logical partition thereof. It could also be a path transparently cross-connected across the access network, such as a PVP in an ATM network.

30 The bandwidth provisioned for the path is determined from an expected traffic load. The sum of the bandwidths provisioned for each

virtual connection may exceed the bandwidth provisioned for the path, thereby introducing a non-null probability for a virtual connection not getting the bandwidth it has been provisioned for.

The multiplexing can be manifold:

- 5 - in the second embodiment (claims 3 and 9), the virtual connections are aggregated one with each other over the path, the CAC means being preliminarily disabled so as to aggregate more virtual connections than there are bandwidth available,
- 10 - in the third embodiment (claims 4 and 10), a virtual connection is dynamically connected to the path when the first service to be delivered over that connection is granted, and dynamically disconnected when the last granted service is completed.

In both embodiments, one further identifies the path over which the so-identified virtual connection is multiplexed, next one checks whether
15 the so-identified path can convey the required bandwidth with the required QoS while maintaining the QoS agreed with the subscribers already sharing that path. This additional checking step is a requisite for preserving the integrity of the services delivered over the path.

Another characterizing embodiment of the present invention is
20 defined in the claims 5 and 6. The bandwidth provisioned for the path is determined according to a statistical traffic law, given a number of virtual connections multiplexed over the path, an estimated or measured traffic load per user and a service deny probability. The way around, a path may also get a pre-determined bandwidth, such as the one enforced by a physical
25 link. Then, the number of virtual connections that are multiplexed over the path is determined according to a traffic law, given that pre-determined available bandwidth, an estimated or measured traffic load per user and a service deny probability. The traffic law could be for instance the Erlang law, assuming a Poisson distribution of the network usage over the
30 subscribers.

It is to be noticed that the term 'comprising', used in the claims, should not be interpreted as being restricted to the means listed thereafter. Thus, the scope of the expression 'a device comprising means A and B' should not be limited to devices consisting only of components A and B. It means that
5 with respect to the present invention, the only relevant components of the device are A and B.

Similarly, it is to be noticed that the term 'coupled', also used in the claims, should not be interpreted as being restricted to direct connections only. Thus, the scope of the expression 'a device A coupled to a device B' should not
10 be limited to devices or systems wherein an output of device A is directly connected to an input of device B. It means that there exists a path between an output of A and an input of B which may be a path including other devices or means.

The above and other objects and features of the invention will
15 become more apparent and the invention itself will be best understood by referring to the following description of an embodiment taken in conjunction with the accompanying drawings wherein:

- fig. 1 represents an access network ACCESS according to the invention,
- fig. 2 represents the virtual connections provisioned across the access
20 network ACCESS to deliver video on demand with a guaranteed QoS.

In a preferred embodiment of the present invention, the access network ACCESS is an ATM based data communication network that connects a plurality of subscribers USR1 to USRN to a core network CORE. The core network CORE is part of an IP based data communication
25 network, such as the Internet.

The access network ACCESS comprises the following distinctive network elements:

- an access node ANOD,
- an edge node ENOD.

30 The access network ACCESS further comprises:

- an access resource controller ARC,

- a network manager NM.

The subscribers USR1 to USRN are coupled to the access node ANOD. At respective ones of the subscriber premises, for instance at USR_i with *i* being an integer from 1 to N, a DTE TERM_i, such as a personal
5 computer, is coupled to a DCE MODEM_i, such as a Digital Subscriber Line (DSL) modem. The DCE MODEM_i is coupled via a transmission line, such as an unshielded twisted copper conductor pair, to the access node ANOD.

The access node ANOD is coupled to the edge nodes ENOD. The edge nodes ENOD is coupled to the core network CORE, so is the access
10 resource controller ARC. The network manager NM is coupled to network elements of the access network ACCESS, including – but not restricted to – the network elements ANOD and ENOD.

The access node ANOD accommodates the protocol suites applicable to the traffic planes, the signalling planes and the management
15 planes terminating at or relayed across any of its physical ports, as well as all the necessary hardware components for those planes to operate.

More specifically and with respect to the scope of the present invention, the access node ANOD accommodates a plurality of physical Input/Output (I/O) ports adapted to terminate an electrical signal of a
20 subscriber line, such as a DSL signal, and a plurality of physical I/O ports adapted to terminate an electrical or optical signal from a network element of the access network ACCESS, such as an optical STM-1 signal. The access node ANOD comprises an ATM switch fabric to switch ATM cells from any incoming ATM virtual channel to any outgoing ATM virtual channel. The
25 access node ANOD accommodates the necessary queuing means on the ATM egress ports. To meet ATM QoS requirements, the access node ANOD implements a suitable cell scheduling policy in association with ATM traffic management functions, including – but not restricted to – CAC, UPC, NPC and CLP. The access node ANOD houses an agent to carry out management
30 operations upon request from the network manager NM. As an example, the access node may be requested to provision a PVC with a given bandwidth

and a given QoS on any ATM port. None of these components is shown in fig. 1 for reasons of readability.

The access node ANOD could be for instance a Digital Subscriber Line Access Multiplexer (DSLAM).

5 The edge node ENOD accommodates the protocol suites applicable to the traffic planes, the signalling planes and the management planes terminating at or relayed across any of its physical ports, as well as all the necessary hardware components for those planes to operate.

10 More specifically, the edge node ENOD accommodates a plurality of physical I/O ports adapted to terminate an electrical or optical signal from a network element of the access network ACCESS, and a plurality of physical I/O ports adapted to terminate an electrical or optical signal from a network element of the core network CORE, such as a 1 Gigabit Ethernet signal. The edge node ENOD is adapted to extract the IP traffic from an
15 incoming ATM virtual channel and to forward it to the core network CORE, and vice-versa. Traffic forwarding is achieved by maintaining on a per subscriber session basis a cross-reference table associating the identifier of an ATM virtual channel with the IP address of the subscriber to which that virtual channel relates. Traffic forwarding could possibly rely upon Protocol
20 Control Information (PCI) of the data packets. The edge node ENOD accommodates the necessary queuing means on the ATM egress ports. To meet ATM QoS requirements, the edge node ENOD implements a suitable cell scheduling policy in association with ATM traffic management functions, including – but not restricted to – CAC, NPC and CLP. The edge
25 node ENOD is adapted to terminate a subscriber point-to-point dial-in session, such as a PPPoE session, which is established across the access network ACCESS over an ATM virtual channel. The edge node ENOD accommodates Authentication Authorization and Accounting (AAA) means, directly or through a dedicated server. The edge node ENOD may further
30 accommodate firewall means, caching means, etc. The edge node ENOD houses an agent to carry out management operations upon request from the

network manager NM. None of these components is shown in fig. 1 for reasons of readability.

The edge node ENOD could be for instance a Broadband Access Server (BAS).

5 The access resource controller ARC is a computing station comprising a Pentium III processor supplied by Intel Corp, as well as all the necessary hardware as known to the skilled person, such as random access memory, a hard disk, an interruption controller, a DMA controller, a graphic card, I/O peripherals.

10 The access resource controller ARC implements a Bandwidth on Demand (BoD) interface across which a service provider requests for a subscriber a pre-determined bandwidth across the access network ACCESS with a pre-determined QoS. The required bandwidth is granted for a period of time, the value of which being specified by the service provider. The
15 required bandwidth and the required QoS are both specified with respect to the IP data communication network. The required bandwidth is defined as the upper bound of the IP throughput required to deliver the service to the subscriber. The required QoS is defined as the IP class of service to which the service relates, such as the differentiated service code point. Quantified
20 objectives may further be specified for the packet loss ratio, the packet delay and the packet jitter.

The access resource controller ARC is adapted to:

- communicate via the core network CORE with service providers, by implementing the appropriate IP protocol suite,
- 25 - retrieve the configuration data of the access network ACCESS,
- identify the resources of the access network ACCESS that have been provisioned for a class of service and for a subscriber,
- maintain bandwidth counter on a per ATM connection basis, each bandwidth counter holding the value of the bandwidth that remains
30 available for a connection,

- check whether an ATM connection can convey a required bandwidth with a required QoS.

The network manager NM of the access network ACCESS is a computing station, such as one of the above mentioned type.

5 The network manager NM accommodates the protocol suite applicable on the interface between the network manager NM and the network agents, such as the Q.3 protocol suite. The network manager NM makes use of a Management Information Base (MIB) abstracting the resources of the access network ACCESS into managed objects and defining
10 the management operations those managed objects may support. Dealing with the present invention, the network manager NM is adapted to provision a given bandwidth across the access network ACCESS with a given QoS, being a PVC or a PVP, by issuing the appropriate management commands to the appropriate agents.

15 The subscribers USR1 to USRN are committed a service level agreement, which includes a guaranteed IP throughput for streaming application such as radio or video broadcasting. A service deny probability is mentioned therein, notifying the subscriber that he might not get the committed bandwidth with a pre-determined probability.

20 Thereupon, PVCs with suitable ATM service categories are configured between the subscribers USR1 to USRN and the edge node ENOD. In the preferred embodiment, each subscriber is supplied with a UBR PVC for best effort traffic and a CBR PVC for real-time traffic. The PCR of the CBR PVCs is determined from the committed IP throughput by
25 including the overhead introduced by the access network ACCESS, such as PPPoE overhead, and by translating the resulting throughput into ATM cell rate.

 The forwarding of the traffic onto the right PVC is out of the scope of the present invention. An algorithm based upon the class of service to
30 which the data packets relate is one possible solution, out of a number of valid alternatives as known to the skilled person. The traffic bifurcation is

performed in the subscriber's DCE MODEMi for the upstream traffic, and in the edge node ENOD for the downstream traffic.

The network manager NM requests the agents of the access node ANOD and the edge node ENOD to:

- 5 - provision N UBR PVCs U11 to U1N between respective ones of the subscriber's DCEs MODEM1 to MODEMN and the access node ANOD,
- provision N UBR PVCs U21 to U2N between the access node ANOD and the edge node ENOD,
- connect the N PVCs U11 to U1N with respective ones of the N PVCs U21
- 10 to U2N, thereby providing N end-to-end UBR connections U1 to UN between respective ones of the subscribers USR1 to USRN and the edge node ENOD,
- provision N CBR PVCs C11 to C1N between respective ones of the subscriber's DCEs MODEM1 to MODEMN and the access node ANOD,
- 15 - provision N CBR PVCs C21 to C2N between the access node ANOD and the edge node ENOD,
- connect the N PVCs C11 to C1N with respective ones of the N PVCs C21 to C2N, thereby providing N end-to-end CBR connections C1 to CN between respective ones of the subscribers USR1 to USRN and the edge
- 20 node ENOD,
- provision a CBR PVP C between the access node ANOD and the edge node ENOD, the bandwidth of which being determined from the real-time traffic expected for the subscribers USR1 to USRN,
- disable the CAC means on the ports of the access node ANOD where the
- 25 PVP C has been provisioned,
- disable the CAC means on the ports of the edge node ENOD where the PVP C has been provisioned,
- aggregate the N PVCs C1 to CN over the PVP C.

Only a few connections are shown in the fig. 2 in order to not
30 overload the drawing.

The PCR to be provisioned for the PVP C is derived from the Erlang B formula. This formula is used to calculate the Flow Blocking Probability (*FBP*) of a multiplexing device:

$$FBP(n, \rho, m) = \frac{(n \cdot \rho)^m / m!}{\sum_{i=0}^m (n \cdot \rho)^i / i!}$$

5 In this formula, n denotes the number of inputs of said multiplexing device, m denotes the number of outputs and ρ denotes the mean load at the n inputs. ρ is comprised between 0 (an input is never used) and 1 (an input is always used). The multiplexing ratio of said multiplexing device is n/m .

10 The value of m is calculated with fixed values of n and ρ and a desired *FBP*. A typical value for the desired *FBP* is 10^{-4} , a typical value for ρ is 10^{-1} . The calculation process makes use of an iterative process, calculating the value of *FBP* for increasing values of m , until the value for *FBP* obtained using the Erlang B formula is smaller than the desired *FBP*.

15 The aggregation of the PVCs C1 to CN over the PVP C can be assimilated to such a multiplexing device wherein:

- $n = N$,
- *FBP* is the above mentioned service deny probability,
- the PCR of the PVP C is m times the PCR of the PVC Ci.

20 Following the request from the network manager NM, the agents of the access node ANOD and the edge node ENOD carry out resource assignment and connection activation. At that stage, the pre-configuration of the access network ACCESS is completed in conformance with the present invention.

25 The configuration data of the access network ACCESS are encoded in a file, which is subsequently loaded into the access resource controller ARC by means of a floppy disk.

In the following, the operation of the preferred embodiment will be described by reference to a service provider VOD coupled to the core

network CORE and delivering to the subscribers USR1 to USRN videos on demand.

A subscriber USR_i requests a video V to the service provider VOD via the UBR connection U_i and next via the core network CORE.

5 The video V requires a bandwidth B_{IP} with a QoS Q_{IP} for a period of time T. The service provider VOD requests the access resource controller ARC via the core network CORE whether it can be granted those resources for the subscriber USR_i.

10 First of all, the access resource controller ARC translates the IP data transfer capabilities as required by the service provider VOD into ATM data transfer capabilities. B_{IP} is translated into B_{ATM} by including the overhead introduced by the access network ACCESS and by translating the resulting throughput into ATM cell rate. Q_{IP} is translated into Q_{ATM} by matching the IP class of service required by the service provider with a
15 suitable ATM service category, possibly assigning values to the cell loss, cell delay and cell jitter objectives.

 The access resource controller ARC retrieves the configuration of the access network ACCESS and identifies the PVC C_i as being capable of conveying the required bandwidth B_{ATM} with the required QoS Q_{ATM}
20 between the edge node ENOD and the subscriber USR_i. Next, the access resource controller ARC looks for any contention point in the access network that may narrow the original bandwidth. The access resource controller ARC identifies the PVP C as such.

 A set of counters hold the value of the bandwidth that remains
25 available for the PVCs C₁ to C_N and for the PVP C. These counters are decremented whenever bandwidth is granted to a service provider and incremented when the resources are released.

 The access resource controller ARC checks whether the counter related to the PVC C_i and the PVP C are both greater or equal than B_{ATM}.
30 If so, the bandwidth is granted and both counters are decremented

accordingly. A timer is set to T seconds and started therefrom. Upon timer expiry, both counter are incremented back to their original value.

Eventually, the access resource controller ARC grants or denies the required bandwidth to the service provider VOD via the core network
5 CORE, which in turns grants or denies the video V to the subscriber USR_i via the core network CORE and next via the UBR connection U_i.

If the service is granted, the ordered video V starts being broadcasted to the subscriber USR_i. The video data packets are flagged with a suitable class of service value and are transmitted over the core network
10 CORE. The data packets are routed through the core network CORE up to the edge node ENOD, wherein they are encapsulated into ATM cells and delivered over the CBR connection C_i to the subscriber USR_i.

In the preferred embodiment, the same bandwidth is provisioned for the upstream and downstream traffic. ATM support asymmetric
15 connection wherein the PCR specified for the egress direction is distinct from the PCR specified for the ingress direction. Such a feature could be helpful for an Asymmetric Digital Subscriber Line (ADSL) access wherein the upstream traffic is allocated a lower bandwidth on the subscriber line portion than the downstream traffic is. The reservation and the supervision
20 of the resources would then be instantiated on a per traffic flow direction basis.

In an alternative embodiment of the present invention, the access resource controller ARC would be coupled to the network manager NM to retrieve the configuration of the access network ACCESS. The access
25 resource controller ARC may load the configuration of the access network ACCESS from the network manager NM via a M3 interface, as defined in the document entitled '*Customer Network Management for ATM Public network Service (M3 Specification)*', published by the ATM forum technical committee on October 1994, revision 1.04, reference AF-NM-0019.000. The
30 access resource controller could also dump a file, e.g. through an FTAM

session, wherein the configuration of the access network ACCESS has been encoded according to a proprietary-format.

In an alternative embodiment of the present invention, the access resource controller ARC would be coupled to the access node ANOD and to the edge node ENOD to read local performance counters. The access node ANOD and the edge node ENOD would implement counters measuring the occupied bandwidth on a per ATM connection basis. Upon bandwidth request, the access resource controller ARC would check whether a PVC or a PVP can convey the required bandwidth by reading the real-time value of the corresponding counters and by checking whether the PVC or the PVP can further accommodate the additional bandwidth. In so doing, the access resource controller does not need to maintain any counter or any timer on its own.

In an alternative embodiment of the present invention, PVCs with other ATM service categories would be provisioned, such as rt-VBR or ABR service category. The bandwidth required via the BoD interface would further include additional parameters along with the maximum IP throughput, such as the mean IP throughput or the minimum required IP throughput. In so doing, the access resource controller ARC would reserve the network resources more efficiently by relying on the statistical averaging of non-correlated traffic flows aggregated to each other.

A person skilled in the art could consider any other label switched technology than ATM, such as MPLS or frame relay technology, as a valid alternative for the transport layer in the access network ACCESS.

A final remark is that embodiments of the present invention are described above in terms of functional blocks. From the functional description of these blocks, given above, it will be apparent for a person skilled in the art of designing electronic devices how embodiments of these blocks can be manufactured with well-known electronic components. A detailed architecture of the contents of the functional blocks hence is not given.

While the principles of the invention have been described above in connection with specific apparatus, it is to be clearly understood that this description is made only by way of example and not as a limitation on the scope of the invention, as defined in the appended claims.